TRUSTED FUTURE

Comments to the Office of the Privacy Commissioner of Canada
on the Draft Guidance *Assessing whether and how to use age assurance – Guidance for websites and online services*

3 August 2026

At Trusted Future, we believe that the connected digital world provides important opportunities for children to learn, grow, and connect in ways that were unimaginable just a generation ago. But there are also new dangers and risks for children in this interconnected, digital world. If we cannot trust that children will be kept safe online, then we put at risk the incredible benefits that come with a digitally connected world. We believe in approaching this issue holistically, not looking for the one simple solution or silver bullet, but seeking to find ways for parents, innovators and policymakers to work together to protect kids and teens online.

It is with this belief that we offer these comments to the Office of the Privacy Commissioner (OPC) draft guidance on, *Assessing whether and how to use age assurance – Guidance for websites and online services* (the "Guidance"), published for consultation on May 4, 2026. We commend the OPC for attempting to create a risk-based, proportionality-driven framework rather than treating age assurance as a default gate to the Internet, and for anchoring the Guidance in the principle that age assurance should be only one of several tools used to protect children online, not a blanket precondition for access.

Our comments draw on our policy analysis of similar measures in the U.S., Europe, and Australia; original survey research we have conducted; and decades of experience working on these issues.

Overview

We appreciate the Guidance's core principles that organizations must first determine if there is a need to differentiate users by age (Section 2), then select a proportionate method calibrated to the risk being addressed (Section 3), and finally implement that method in a privacy-protective way that avoids secondary uses, profiling, and unnecessary repetition (Section 4). This structure correctly rejects "silver bullet" approaches to age assurance in favor of a harm-specific analysis that is consistent with the limitations of age assurance technologies.

As our research has found that no age assurance technology available today is perfect or free of significant trade-offs, and policy approaches that assume otherwise tend to generate new privacy and equity risks without reliably achieving the safety outcomes they promise.

We also support the Guidance's emphasis on proportionate design and its clear prohibitions on secondary use. These principles align closely with the age-assurance frameworks that we believe work best in practice: those that rely on privacy-preserving age ranges rather than exact birthdates or verified identity, that minimize the number of parties that receive a child's age signal, and that avoid creating centralized repositories of sensitive identity or biometric data.

With that said, we believe the final guidance would benefit from four points discussed below:

(A) More clearly recognizing privacy risks and challenges of existing assurance methods;

(B) Cautioning specifically against facial-recognition-based age estimation given documented accuracy disparities;

(C) Avoiding ecosystem-wide, indiscriminate collection of verified age data; and

(D) Highlighting age-range signals as a useful privacy-protective default.

Specific Comments and Recommendations

(A) More clearly recognize privacy risks and challenges of existing assurance methods

Section 3.1 correctly ties the necessary "level of effectiveness" of an age assurance method to the severity of the underlying risk and acknowledges that age estimation based on biometric information or verification based on government ID may be appropriate only where a high level of effectiveness is warranted and the method is implemented in a privacy-protective manner.

We appreciate that OPC's Privacy and age assurance – Exploratory consultation identified a "Trust Deficit" as a barrier for consumer and correctly pointed out that, "Even if privacy-enhancing technologies are implemented and age assurance services are subject to strict oversight, it is likely that some individuals will forego accessing content or services that requires age assurance or seek that content in riskier ways (such as via the dark web)."

We recommend the final guidance more fully recognizes the privacy and equity costs of the main categories of age assurance methods, so that relying parties conducting the necessity/proportionality analysis in Section 3.1 have a clearer basis for comparison.

In our review of the available approaches, the trade-offs are significant:

- *Self-attestation* is low-friction and minimally invasive, but it depends on the user being honest and is easily circumvented. For example, in the U.S. nearly 40% of children ages 8–12 use social media according to the Surgeon General, however,

according to the Children's Online Privacy Protection Act (COPPA) age 13 is the minimum age permitted by law, demonstrating a wide gap between restrictions on paper and the reality in practice. This method should not be relied upon where the potential harm is significant.

- *Behavioral age estimation*, which infers age from usage patterns, content, or account signals avoids an explicit verification step, but requires organizations to collect and analyze substantially more personal information about all users in order to generate the inference. This creates exactly the kind of profiling risk the guidance seeks to avoid. As the OPC's survey from May of this year found, youth participants have a low knowledge base about what companies collect and retain about them, and about their privacy rights and how to exercise them.
- *Facial-recognition-based age estimation* raises accuracy and discrimination concerns that we discuss separately below.
- *Government-ID verification* is the most reliable method but also has two major risks: it is the highest-risk from a data protection standpoint and the most exclusionary. About two thirds of Canadians are not comfortable having face or photo scanned to verify age online. This concern is well founded as government-ID and other document-based verification methods also create concentrated repositories of highly sensitive personal information that can be the target of data breaches. Meanwhile, only about 70% of Canadians have a valid drivers license. Those without a government issued ID often come from categories of underrepresented populations including those experiencing homelessness, indigenous communities, low-income, and senior citizens. Such a method would unfairly deny access to these legitimate users.
- *Credit Card verification* is another proposal. But here again, it is exclusionary. Around 11% to 15% of adult Canadians do not have a credit card.

It is possible that in the future age verification and age assurance systems will be able to avoid the worst problems of the current systems. But for the time being, standards-based solutions (like website metadata) combined with some of the smart parental control tools that allow parents to block and filter mature content, and to individually approve which websites their children can access, provide the most practical opportunities both to protect children from adult content and to protect the rights of adults to access the content they want, while also avoiding severe privacy and security issues.

(B) Caution specifically against facial-recognition-based age estimation given documented accuracy disparities

Facial-recognition-based age estimation is increasingly offered by third-party age-assurance vendors as a "less intrusive" alternative to document verification. We believe the final guidance should specifically address its accuracy limitations which are not evenly distributed across the population and therefore raise equity concerns distinct from the general privacy concerns addressed elsewhere in the Guidance.

Independent evaluation by the U.S. National Institute of Standards and Technology (NIST) of several commercial age-estimation providers found average estimation errors of roughly three years, with error rates varying significantly by skin color, gender, and region of birth. Just 7.2% to 34.5% of 13-year-olds were correctly estimated within one year of their actual age. For this and other reasons, the U.S. Federal Trade Commission (FTC) declined to authorize the use of facial-recognition age-estimation method to satisfy verifiable-parental-consent requirements under COPPA (Children's Online Privacy Protection Act) in 2024.

We recommend the final guidance direct organizations to treat facial-recognition-based age estimation as high-risk from an accuracy and non-discrimination standpoint, to disclose known error rates and demographic disparities to users, and to provide a meaningful, low-friction appeal path for individuals incorrectly flagged by such methods.

(C) Avoid ecosystem-wide, indiscriminate collection of verified age data

Section 2.2.2 appropriately limits the requirement to use age assurance to organizations whose site or service are likely to be accessed by a non-trivial number of children and whose personal information practices pose a potential harm to those children.

To the extent that OPC is defining “online services” as including mobile apps, then it needs to avoid a mobile ecosystem-wide collection of verified age data. We caution against legislative or regulatory proposals that would require a device, operating system, browser, or app marketplace to verify and transmit every user's age or a specific age credential to every downloading application, or website regardless of whether that particular application's content or practices pose any potential harm to children. We recommend the final guidance specifically reject such a methodology.

We raise this because several proposals that would require mobile operating systems to broadcast this kind of sensitive information to every app are currently under active consideration in the United States and we believe Canada should avoid following this mistaken path. Under this app-based model, a general-purpose application with no age-inappropriate content or practice, such as a weather app, a news app, a local retailer's app, would still receive a verified age signal for every user, including adult users, as a condition of a single download. This is difficult to reconcile with the Guidance's own proportionality principle laid out in Section 3.1, which requires that individuals "not be required to provide more, or more sensitive, personal information than is necessary to adequately address the identified risk(s)". An application with no age-related risk has, by definition, no risk that a verified age signal is necessary to address. It is also difficult to reconcile with Section 2.2.2's non-trivial-access threshold, since it would apply age verification requirements to millions of applications regardless of whether children are likely to access them at all.

This model also has the practical effect of multiplying, rather than minimizing, the number of parties that hold sensitive age or identity data — the opposite of the data-minimization

objective in the Guidance. These proposals distribute a verified age credential to every application developer at the point of download rather than limiting collection to the smaller number of services whose content or practices actually warrant differentiation by age. This creates a correspondingly larger number of independent repositories of sensitive data, each a potential point of failure for a privacy violation or data breach.

We would also note, as a practical matter, that a device- or app-store-level check applied only at the point of download is easy to work around. A user who cannot access age-restricted material through a blocked application can often still reach the same content through a mobile browser or desktop computer, meaning this model does not deliver the sustained protection the guidance's own risk-based framework is designed to achieve, and presents easy workarounds for tech-savvy children.

Furthermore, we have not yet seen that broad based age-restrictions have been effective in areas where they have been tried. The clearest example is in Australia, where a ban on children and teens under 16 took effect last December. We are now getting helpful data on the effectiveness of the ban. According to a new study published in the British Medical Journal, more than 85% of kids and teens under 16 who participated reported continuing to use social media after the ban took effect. The eSafety Commission's own survey found that nearly 70% of under-16s with accounts on Instagram, Snapchat or TikTok had maintained access. The eSafety Commission says the technology being used by the companies, like facial age estimation, was not effective enough and argued the firms had lax guardrails. But these results also suggest that workarounds from kids and teens, such as VPNs, have been widespread.

Even in the United Kingdom, where the Online Safety Act's child safety duties came into force a year ago, the Department for Science, Innovation and Technology (DSIT) recently released a report that found 26% of 11- to 17-year-olds use a VPN, though the most common reason cited was for privacy.

These bans have not only been ineffective, they have also come at a cost, delivering negative unintended consequences for many kids. For example, in Australia, more than half of under 16 significantly impacted by the ban are now getting less news according to a recent survey. There are also reports of children with disabilities becoming more and more isolated without the connections that they used to reach online.

We believe this clarification is consistent with, rather than in tension with, the existing text of the Guidance. It would simply pre-empt an interpretation that could otherwise be used to justify precisely the kind of disproportionate, unnecessary collection the guidance is designed to prevent.

(D) Highlight useful existing privacy-protective defaults

There already exist some powerful, privacy-protective parental controls that some companies offer today that put the parents in control while minimizing the amount of data that is collected and stored. For example, Apple's Safari browser includes several built-in

tools designed to protect children online. Content restrictions let parents limit access to adult websites, either by blocking explicit content automatically or allowing only a specific list of approved sites. Apple also recently introduced “Ask to Browse,” which requires kids to get a parent's permission before visiting a new website for the first time, working across iPhone, iPad, and Mac. This builds on the existing Child Account system, which is required for children under 13 and available up to age 18, and which activates a layered set of age-appropriate protections across the device, including content filtering in Safari. Parents can also manage these settings remotely and receive alerts.

To the extent that OPC is defining “online services” as including mobile apps, we would note that both Google and Apple, the two major mobile operating systems, now offer parent-managed, on-device age-range signals. These are tools that let a parent establish a child's age range once, at the device or account level, and share only that range with participating applications with parental consent, rather than requiring each site or app to separately verify a birthdate or identity document. Complementary on-device tools such as content-download approval, communication-permission controls, on-device detection and blurring of explicit imagery, and app-level developer APIs that let a participating application request only an age-range signal rather than a specific date of birth, are already deployed at scale.

These mechanisms illustrate the kind of privacy-protective design the guidance is describing by using the least sensitive category of information necessary, such as a range rather than a birthdate or identity document; keep the parent in control of when and with whom the signal is shared; avoid creating new centralized repositories of identity documents; and allow the signal to be reused across services rather than re-collected at every destination.

Conclusion

We appreciate the OPC's overall approach in this Guidance, pursuing harm-specific, and proportionate framework that treats age assurance as one tool among several, applies it only where genuinely necessary, and requires that it be designed and operated in a privacy-protective manner. But on an issue as important and complex as age assurance, we believe the important details and considerations outlined above are important for the OPC to take into account as it finalizes the Guidance.

We would welcome the opportunity to discuss these comments further with OPC staff and thank the Office for the opportunity to submit these comments.